

Chapter 8

Layer 3 VPN Configuration Troubleshooting Guidelines

This chapter discusses the following strategies and tools for troubleshooting Layer 3 virtual private network (VPN) configurations:

Diagnose Common Problems on page 99

Use the Ping and Traceroute Commands to Troubleshoot Layer 3 VPN Topologies on page 103

Indirect Next-hop Address Space and Route Reflectors on page 110

Diagnose Common Problems

When problems arise in a Layer 3 VPN configuration, the best way to troubleshoot is to start at one end of the VPN (the local customer edge [CE] router) and follow the routes to the other end of the VPN (the remote CE router). The following troubleshooting steps should help you diagnose common problems:

1. If you configured a routing protocol between the local provider edge (PE) and CE routers, check that the peering and adjacency are fully operational. When you do this, be sure to specify the name of the routing instance. For example, to check OSPF adjacencies, enter the command `show ospf neighbor instance routing-instance-name` on the PE router.

If the peering and adjacency are not fully operational, check the routing protocol configuration on the CE router and check the routing protocol configuration for the associated VPN routing instance on the PE router.

2. Check that the local CE and PE routers can ping each other.

To check that the local CE router can ping the VPN interface on the local PE router, use a ping command in the following format, specifying the IP address or name of the PE router:

```
ping (ip-address | host-name)
```

To check that the local PE router can ping the CE router, use a ping command in the following format, specifying the IP address or name of the CE router, the name of the interface used for the VPN, and the source IP address (the local address) in outgoing ECHO_REQUEST packets:

```
ping ip-address vpn-interface interface local echo-address
```

Often, the peering or adjacency between the local CE and local PE routers needs to come up before a ping command is successful. To check that a link is operational in a lab setting, remove the interface from the VRF by deleting the interface statement from the [edit routing-instance *routing-instance-name*] hierarchy level and recommitting the configuration. Doing this removes the interface from the VPN. Then try the ping command again. If the command is successful, configure the interface back into the VPN and check the routing protocol configuration on the local CE and PE routers again.

3. On the local PE router, check that the routes from the local CE router are in the VRF routing table (*routing-instance-name*.inet.0):

```
show route table routing-instance-name.inet.0 [detail]
```

The following example shows the routing table entries. Here, the loopback address of the CE router is 10.255.14.155/32 and the routing protocol between the PE and CE routers is BGP. The entry looks like any ordinary BGP announcement.

```
10.255.14.155/32 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Nexthop: 192.168.197.141 via fe-1/0/0.0, selected
            State: <Active Ext>
            Peer AS:      1
            Age: 45:46
            Task: BGP_1.192.168.197.141+179
            Announcement bits (2): 0-BGP.0.0.0.0+179 1-KRT
            AS path: 1 I
            Localpref: 100
            Router ID: 10.255.14.155
```

If the routes from the local CE router are not present in the VRF routing table, check that the CE router is advertising routes to the PE router. If static routing is used between the CE and PE routers, make sure the proper static routes are configured.

4. On a remote PE router, check that the routes from the local CE router are present in the bgp.l3vpn.0 routing table:

```
show route table bgp.l3vpn.0 extensive
```

The following example shows the routing table entries.

```
10.255.14.175:3:10.255.14.155/32 (1 entry, 0 announced)
  *BGP      Preference: 170/-101
            Route Distinguisher: 10.255.14.175:3
            Source: 10.255.14.175
            Nexthop: 192.168.192.1 via fe-1/1/2.0, selected
            label-switched-path vpn07-vpn05
            Push 100004, Push 100005(top)
            State: <Active Int Ext>
            Local AS:      69 Peer AS:      69
            Age: 15:27      Metric2: 338
            Task: BGP_69.10.255.14.175+179
            AS path: 1 I
            Communities: target:69:100
            BGP next hop: 10.255.14.175
            Localpref: 100
            Router ID: 10.255.14.175
            Secondary tables: VPN-A.inet.0
```

The output of the `show route table bgp.l3vpn.0` extensive command contains the following information specific to the VPN:

In the prefix name (the first line of the output), the route distinguisher is added to the route prefix of the local CE router. Because the route distinguisher is unique within the Internet, the concatenation of the route distinguisher and IP prefix provides unique VPN-IPv4 routing entries.

The Route Distinguisher field lists the route distinguisher separately from the VPN-IPv4 address.

The label-switched-path field shows the name of the LSP used to carry the VPN traffic.

The Push field shows both labels being carried in the VPN-IPv4 packet. The first label is the inner label, which is the VPN label that was assigned by the PE router. The second label is the outer label, which is an RSVP label.

The Communities field lists the target community.

The Secondary tables field lists other routing tables on this router into which this route has been installed.

If routes from the local CE router are not present in the `bgp.l3vpn.0` routing table on the remote PE router, do the following:

Check the VRF import filter on the remote PE router, which is configured in the `vrf-import` statement. (On the local PE router, you check the VRF export filter, which is configured with the `vrf-export` statement.)

Check that there is an operational LSP or an LDP path between the PE routers. To do this, check that the IBGP next-hop addresses are in the `inet.3` table.

Check that the IBGP session between the PE routers is established and configured properly.

Check for “hidden” routes, which usually means that routes were not labeled properly. To do this, use the `show route table bgp.l3vpn.0 hidden` command.

Check that the inner label matches the inner VPN label that is assigned by the local PE router. To do this, use the `show route table mpls` command.

The following example shows the output of this command on the remote PE router. Here, the inner label is 100004.

```
...
Push 100004, Push 10005 (top)
```

The following example shows the output of this command on the local PE router, which shows that the inner label of 100004 matches the inner label on the remote PE router:

```
...
100004          *[VPN/7] 06:56:25, metric 1
                > to 192.168.197.141 via fe-1/0/0.0, Pop
```

5. On the remote PE router, check that the routes from the local CE router are present in the VRF table (*routing-instance-name.inet.0*):

```
show route table routing-instance-name.inet.0 [detail]
```

The following example shows the routing table entries.

```
10.255.14.155/32 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Route Distinguisher: 10.255.14.175:3
            Source: 10.255.14.175
            Nexthop: 192.168.192.1 via fe-1/1/2.0, selected
            label-switched-path vpn07-vpn05
            Push 100004, Push 100005(top)
            State: <Secondary Active Int Ext>
            Local AS: 69 Peer AS: 69
            Age: 1:16:22 Metric2: 338
            Task: BGP_69.10.255.14.175+179
            Announcement bits (2): 1-KRT 2-VPN-A-RIP
            AS path: 1 I
            Communities: target:69:100
            BGP next hop: 10.255.14.175
            Localpref: 100
            Router ID: 10.255.14.175
            Primary Routing Table bgp.l3vpn.0
```

In this routing table, the route distinguisher is no longer prepended to the prefix. The last line, Primary Routing Table, lists the table from which this route was learned.

If the routes are not present in this routing table, but were present in Step 4, the routes might have not passed the VRF import policy on the remote PE router.

If a VPN-IPv4 route matches no vrf-import policy, the route does not show up in the bgp.l3vpn table at all and hence is not present in the VRF table. If this occurs, it might indicate that on the PE router, you have configured another vrf-import statement on another VPN (with a common target), and the routes show up in the bgp.l3vpn.0 table, but are imported into the wrong VPN.

6. On the remote CE router, check that the routes from the local CE router are present in the routing table (inet.0):

```
show route
```


If the routes are not present, check the routing protocol configuration between the remote PE and CE routers, and make sure that peers and adjacencies (or static routes) between the PE and CE routers are correct.

7. If, in Steps 1 through 6, you have determined that routes originated from the local CE router are correct, check the routes originated from the remote CE router by repeating Steps 1 through 6.

Use the Ping and Traceroute Commands to Troubleshoot Layer 3 VPN Topologies

This section provides examples of how to use the ping command to check the accessibility of various routers in a VPN topology, and how to use the traceroute command to check the path that packets travel between the VPN routers. The topology shown in Figure 15 illustrates these commands.

Figure 15: Layer 3 VPN Topology for Ping and Traceroute Command Examples



Ping One CE Router from the Other

You can ping one CE router from the other by specifying the other CE router's loopback address as the IP address in the ping command. This ping command succeeds if the loopback addresses have been announced by the CE routers to their directly connected PE routers. The success of these ping commands also means that Router CE1 can ping any network devices beyond Router CE2, and vice versa. See Figure 15 for the topology referenced in these examples.

Ping Router CE2 (VPN5) from Router CE1 (VPN4):

```
user@vpn4> ping 10.255.10.5 local 10.255.10.4 count 3
PING 10.255.10.5 (10.255.10.5): 56 data bytes
64 bytes from 10.255.10.5: icmp_seq=0 ttl=253 time=1.086 ms
64 bytes from 10.255.10.5: icmp_seq=1 ttl=253 time=0.998 ms
64 bytes from 10.255.10.5: icmp_seq=2 ttl=253 time=1.140 ms

--- 10.255.10.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.998/1.075/1.140/0.059 ms
```

To determine the path from Router CE1's loopback interface to Router CE2's loopback interface, use the following traceroute command:

```
user@vpn4> traceroute 10.255.10.5 source 10.255.10.4
traceroute to 10.255.10.5 (10.255.10.5) from 10.255.10.4, 30 hops max, 40 byte
packets
 1  vpn1-fe-110.isp-core.net (192.168.192.1)  0.680 ms  0.491 ms  0.456 ms
 2  vpn2-t3-001.isp-core.net (192.168.192.110) 0.857 ms  0.766 ms  0.754 ms
    MPLS Label=100005 CoS=0 TTL=1 S=1
 3  vpn5.isp-core.net (10.255.10.5)  0.825 ms  0.886 ms  0.732 ms
```

Ping Router CE1 (VPN4) from Router CE2 (VPN5):

```
user@vpn5> ping 10.255.10.4 local 10.255.10.5 count 3
PING 10.255.10.4 (10.255.10.4): 56 data bytes
64 bytes from 10.255.10.4: icmp_seq=0 ttl=253 time=1.042 ms
64 bytes from 10.255.10.4: icmp_seq=1 ttl=253 time=0.998 ms
64 bytes from 10.255.10.4: icmp_seq=2 ttl=253 time=0.954 ms

--- 10.255.10.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.954/0.998/1.042/0.036 ms
```

To determine the path from Router CE2 to Router CE1, use the following traceroute command:

```
user@vpn5> traceroute 10.255.10.4 source 10.255.10.5
traceroute to 10.255.10.4 (10.255.10.4) from 10.255.10.5, 30 hops max, 40 byte
packets
 1  vpn-08-t3-003.isp-core.net (192.168.193.2)  0.686 ms  0.519 ms  0.548 ms
 2  vpn1-so-100.isp-core.net (192.168.192.100) 0.918 ms  0.869 ms  0.859 ms
    MPLS Label=100021 CoS=0 TTL=1 S=1
 3  vpn4.isp-core.net (10.255.10.4)  0.878 ms  0.760 ms  0.739 ms
```

Ping the Remote PE and CE Routers from the Local CE Router

From the local CE router, you can ping the VPN interfaces on the remote PE and CE routers, which are point-to-point interfaces. See Figure 15 for the topology referenced in these examples.

Ping Router CE2 (VPN5) from Router CE1 (VPN4):

```
user@vpn4> ping 192.168.193.5 local 10.255.10.4 count 3
PING 192.168.193.5 (192.168.193.5): 56 data bytes
64 bytes from 192.168.193.5: icmp_seq=0 ttl=253 time=1.040 ms
64 bytes from 192.168.193.5: icmp_seq=1 ttl=253 time=0.891 ms
64 bytes from 192.168.193.5: icmp_seq=2 ttl=253 time=0.944 ms

--- 192.168.193.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.891/0.958/1.040/0.062 ms
```

To determine the path from Router CE1's loopback interface to Router CE2's directly connected interface, use the following traceroute command:

```
serpil@vpn4> traceroute 192.168.193.5 source 10.255.10.4
traceroute to 192.168.193.5 (192.168.193.5) from 10.255.10.4, 30 hops max, 40
byte packets
 1  vpn1-fe-110.isp-core.net (192.168.192.1)  0.669 ms  0.508 ms  0.457 ms
 2  vpn2-t3-001.isp-core.net (192.168.192.110)  0.851 ms  0.769 ms  0.750 ms
    MPLS Label=100000 CoS=0 TTL=1 S=1
 3  vpn5-t3-003.isp-core.net (192.168.193.5)  0.829 ms  0.838 ms  0.731 ms
```

Ping Router PE2 (VPN2) from Router CE1 (VPN4). In this case, packets that originate at Router CE1 go to Router PE2, then to Router CE2, and back to Router PE2 before Router PE2 can respond to ICMP requests. You can verify this using the traceroute command.

```
user@vpn4> ping 192.168.193.2 local 10.255.10.4 count 3
PING 192.168.193.2 (192.168.193.2): 56 data bytes
64 bytes from 192.168.193.2: icmp_seq=0 ttl=254 time=1.080 ms
64 bytes from 192.168.193.2: icmp_seq=1 ttl=254 time=0.967 ms
64 bytes from 192.168.193.2: icmp_seq=2 ttl=254 time=0.983 ms

--- 192.168.193.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.967/1.010/1.080/0.050 ms
```

To determine the path from Router CE1 to Router PE2, use the following traceroute command:

```
user@vpn4> traceroute 192.168.193.2 source 10.255.10.4
traceroute to 192.168.193.2 (192.168.193.2) from 10.255.10.4, 30 hops max, 40
byte packets
 1  vpn1-fe-110.isp-core.net (192.168.192.1)  0.690 ms  0.490 ms  0.458 ms
 2  vpn2-t3-003.isp-core.net (192.168.193.2)  0.846 ms  0.768 ms  0.749 ms
    MPLS Label=100000 CoS=0 TTL=1 S=1
 3  vpn5-t3-003.isp-core.net (192.168.193.5)  0.643 ms  0.703 ms  0.600 ms
 4  vpn-08-t3-003.isp-core.net (192.168.193.2)  0.810 ms  0.739 ms  0.729 ms
```

You cannot ping one CE router from the other if the VPN interface is a multiaccess interface, such as the fe-1/1/2.0 interface on Router CE1. To ping Router CE1 from Router CE2, you must configure a static route on Router PE1 to the VPN interface of Router CE1 that has a next hop pointing to Router CE1 (at the [edit routing-instance *routing-instance-name*] hierarchy level) and this route must be announced from Router PE1 to Router PE2. The following configuration portions illustrate this configuration:

```
[edit]
routing-instances {
  direct-multipoint {
    instance-type vrf;
    interface fe-1/1/0.0;
    route-distinguisher 69:1;
    vrf-import direct-import;
    vrf-export direct-export;
    routing-options {
      static {
        route 192.168.192.4/32 next-hop 192.168.192.4;
      }
    }
  }
  protocols {
    bgp {
      group to-vpn4 {
        peer-as 1;
        neighbor 192.168.192.4;
      }
    }
  }
}
policy-options {
  policy-statement direct-export {
    term a {
      from protocol bgp;
      then {
        community add direct-comm;
        accept;
      }
    }
    term b {
      from {
        protocol static;
        route-filter 192.168.192.4/32 exact;
      }
      then {
        community add direct-comm;
        accept;
      }
    }
    term d {
      then reject;
    }
  }
}
```


Now you can ping Router CE1 from Router CE2:

```
user@vpn5> ping 192.168.192.4 local 10.255.10.5 count 3
PING 192.168.192.4 (192.168.192.4): 56 data bytes
64 bytes from 192.168.192.4: icmp_seq=0 ttl=253 time=1.092 ms
64 bytes from 192.168.192.4: icmp_seq=1 ttl=253 time=1.019 ms
64 bytes from 192.168.192.4: icmp_seq=2 ttl=253 time=1.031 ms

--- 192.168.192.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.019/1.047/1.092/0.032 ms
```

To determine the path between these two interfaces, use the following traceroute command:

```
user@vpn5> traceroute 192.168.192.4 source 10.255.10.5
traceroute to 192.168.192.4 (192.168.192.4) from 10.255.10.5, 30 hops max, 40
byte packets
 1 vpn-08-t3003.isp-core.net (192.168.193.2) 0.678 ms 0.549 ms 0.494 ms
 2 vpn1-so-100.isp-core.net (192.168.192.100) 0.873 ms 0.847 ms 0.844 ms
   MPLS Label=100021 CoS=0 TTL=1 S=1
 3 vpn4-fe-112.isp-core.net (192.168.192.4) 0.825 ms 0.743 ms 0.764 ms
```

Ping the Directly Connected PE and CE Routers from Each Other

From the loopback interfaces on the CE routers, you can ping the VPN interface on the directly connected PE router. See Figure 15 for the topology referenced in these examples.

From the loopback interface on Router CE1 (VPN4), ping the VPN interface, fe-1/1/0.0, on Router PE1:

```
user@vpn4> ping 192.168.192.1 local 10.255.10.4 count 3
PING 192.168.192.1 (192.168.192.1): 56 data bytes
64 bytes from 192.168.192.1: icmp_seq=0 ttl=255 time=0.885 ms
64 bytes from 192.168.192.1: icmp_seq=1 ttl=255 time=0.757 ms
64 bytes from 192.168.192.1: icmp_seq=2 ttl=255 time=0.734 ms

--- 192.168.192.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.734/0.792/0.885/0.066 ms
```

To determine the path from the loopback interface on Router CE1 to the VPN interfaces on Router PE1, use the following traceroute command:

```
user@vpn4> traceroute 192.168.192.1 source 10.255.10.4
traceroute to 192.168.192.1 (192.168.192.1) from 10.255.10.4, 30 hops max, 40
byte packets
 1 vpn1-fe-110.isp-core.net (192.168.192.1) 0.828 ms 0.657 ms 1.972 ms
```

From the loopback interface on Router CE2 (VPN5), ping the VPN interface, t3-0/0/3.0, on Router PE2:

```
user@vpn5> ping 192.168.193.2 local 10.255.10.5 count 3
PING 192.168.193.2 (192.168.193.2): 56 data bytes
64 bytes from 192.168.193.2: icmp_seq=0 ttl=255 time=0.998 ms
64 bytes from 192.168.193.2: icmp_seq=1 ttl=255 time=0.834 ms
64 bytes from 192.168.193.2: icmp_seq=2 ttl=255 time=0.819 ms

--- 192.168.193.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.819/0.884/0.998/0.081 ms
```

To determine the path from the loopback interface on Router CE2 to the VPN interfaces on Router PE2, use the following traceroute command:

```
serpil@vpn5> traceroute 192.168.193.2 source 10.255.10.5
traceroute to 192.168.193.2 (192.168.193.2) from 10.255.10.5, 30 hops max, 40
byte packets
 1  vpn-08-t3003.isp-core.net (192.168.193.2)  0.852 ms  0.670 ms  0.656 ms
```

From the VPN interface on the PE router, you can ping the VPN or loopback interface on the directly connected CE router.

From the VPN interface on Router PE1 (VPN1), ping the VPN interface on Router CE1, fe-1/1/0.0:

```
user@vpn1> ping 192.168.192.4 vpn-interface fe-1/1/0.0 local 192.168.192.1 count
3
PING 192.168.192.4 (192.168.192.4): 56 data bytes
64 bytes from 192.168.192.4: icmp_seq=0 ttl=255 time=0.866 ms
64 bytes from 192.168.192.4: icmp_seq=1 ttl=255 time=0.728 ms
64 bytes from 192.168.192.4: icmp_seq=2 ttl=255 time=0.753 ms

--- 192.168.192.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.728/0.782/0.866/0.060 ms
```

From the VPN interface on Router PE1 (VPN1), ping the loopback interface on Router CE1, 10.255.10.4:

```
user@vpn1> ping 10.255.10.4 vpn-interface fe-1/1/0.0 local 192.168.192.1 count 3
PING 10.255.10.4 (10.255.10.4): 56 data bytes
64 bytes from 10.255.10.4: icmp_seq=0 ttl=255 time=0.838 ms
64 bytes from 10.255.10.4: icmp_seq=1 ttl=255 time=0.760 ms
64 bytes from 10.255.10.4: icmp_seq=2 ttl=255 time=0.771 ms

--- 10.255.10.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.760/0.790/0.838/0.034 ms
```

To determine the path from the VPN interface on Router PE1 to the VPN and loopback interfaces on Router CE1, respectively, use the following traceroute commands:

```
user@vpn1> traceroute 10.255.10.4 vpn-interface fe-1/1/0.0 source 192.168.192.1
traceroute to 10.255.10.4 (10.255.10.4) from 192.168.192.1, 30 hops max, 40 byte
packets
 1  vpn4.isp-core.net (10.255.10.4)  0.842 ms  0.659 ms  0.621 ms

user@vpn1> traceroute 192.168.192.4 vpn-interface fe-1/1/0.0 source
192.168.192.1
traceroute to 192.168.192.4 (192.168.192.4) from 192.168.192.1, 30 hops max, 40
byte packets
 1  vpn4-fe-112.isp-core.net (192.168.192.4)  0.810 ms  0.662 ms  0.640 ms
```

From the VPN interface on Router PE2 (VPN2), ping the VPN interface on Router CE2, t3-0/0/3.0:

```
user@vpn2> ping 192.168.193.5 vpn-interface t3-0/0/3.0 local 192.168.193.2
count 3
PING 192.168.193.5 (192.168.193.5): 56 data bytes
64 bytes from 192.168.193.5: icmp_seq=0 ttl=255 time=0.852 ms
64 bytes from 192.168.193.5: icmp_seq=1 ttl=255 time=0.909 ms
64 bytes from 192.168.193.5: icmp_seq=2 ttl=255 time=0.793 ms

--- 192.168.193.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.793/0.851/0.909/0.047 ms
```

From the VPN interface on Router PE2 (VPN2), ping the loopback interface on Router CE2, 10.255.10.5:

```
user@vpn2> ping 10.255.10.5 vpn-interface t3-0/0/3.0 local 192.168.193.2 count 3
PING 10.255.10.5 (10.255.10.5): 56 data bytes
64 bytes from 10.255.10.5: icmp_seq=0 ttl=255 time=0.914 ms
64 bytes from 10.255.10.5: icmp_seq=1 ttl=255 time=0.888 ms
64 bytes from 10.255.10.5: icmp_seq=2 ttl=255 time=1.066 ms

--- 10.255.10.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.888/0.956/1.066/0.079 ms
```

To determine the path from the VPN interface on Router PE2 to the VPN and loopback interfaces on Router CE2, respectively, use the following traceroute commands:

```
user@vpn2> traceroute 10.255.10.5 vpn-interface t3-0/0/3.0 source 192.168.193.2
traceroute to 10.255.10.5 (10.255.10.5) from 192.168.193.2, 30 hops max, 40 byte
packets
 1  vpn5.isp-core.net (10.255.10.5)  1.009 ms  0.677 ms  0.633 ms

user@vpn2> traceroute 192.168.193.5 vpn-interface t3-0/0/3.0 source
192.168.193.2
traceroute to 192.168.193.5 (192.168.193.5) from 192.168.193.2, 30 hops max, 40
byte packets
 1  vpn5-t3-003.isp-core.net (192.168.193.5)  0.974 ms  0.665 ms  0.619 ms
```

Ping a Remote CE Router from a PE Router

If you attempt to ping a remote CE router from a PE router, Internet Control Message Protocol (ICMP) echo requests are sent from the PE router, with the PE router's VPN interface as the source. Other PE routers have a route back to that address with a VPN label. When the echo replies return, they include a label. The PE router pops the VPN label and sends the packet from the VPN interface to the local CE router. The local CE router sends it back to the PE router, its actual destination.

When a Juniper Networks router receives a labeled packet, the label is popped (depending on the label operation specified), and the packet is forwarded to an interface, even if the packet is destined for that particular PE router. Labeled packets are not analyzed further for the IP information under the label.

If there is a problem with the connection to the local CE router, packets are sent out but do not return to the PE router, and the ping fails. So if the connection between your PE router and local CE router is down, sending a ping to the remote CE router fails even though the connection to the remote CE router might be functional.

Disable Normal TTL Decrementing for Layer 3 VPNs

For information on how to disable normal TTL decrementing for Layer 3 VPNs, see “Disable Normal TTL Decrementing for VPNs” on page 22.

Indirect Next-hop Address Space and Route Reflectors

If you attempt to allocate indirect next-hop indexes for all the active routes in the routing table, the indirect next-hop address space can be exhausted. This occurs most frequently in topologies in which a Juniper Networks router acts as a VPN route reflector and a BGP peer with other vendors' routers. The other routers advertise unique label stacks per prefix, instead of sharing common label stacks across many prefixes.

You can configure a forwarding table export policy to prevent routes from being installed in the forwarding table, even when those routes are active in the routing table. With this change, the routes that are omitted from the forwarding table do not have additional indirect next-hop indexes allocated to them.

The following configuration avoids indirect next-hop address space exhaustion, but does not allow the router to forward traffic for the BGP learned prefixes:

```
routing-options {
  forwarding-table {
    export kern;
  }
}
policy-options {
  policy-statement kern {
    from protocol bgp;
    then reject;
  }
}
```